Exhibit K

HWG | HARRIS, WILTSHIRE & GRANNIS LLP

Redacted - For Public Inspection

August 21, 2018

Via Email
Ms. Elizabeth Drogula
Deputy Division Chief
Telecommunications Access Policy Division
Wireline Competition Bureau
Federal Communications Commission
445 12th Street, SW
Washington, DC 20554

Re: GCI Rural Health Care Support for Funding Year 2017

Dear Ms. Drogula,

Pursuant to Sections 0.457 and 0.459 of the Federal Communications Commission's ("FCC") rules, 47 C.F.R. §§ 0.457, 0.459, GCI Communication Corp. ("GCI") hereby requests confidential treatment of the attachments that it submits herein. This submission relates to the Telecommunications Access Policy Division ("the Division") review of outstanding Fiscal Year 2017 funding under the Rural Health Care ("RHC") Telecom Program. Specifically, the attachments include the following (the "Confidential Information"):

- Revised Bandwidth Allocation Methodology for the purposes of allocating costs among services in the TERRA RoR cost study based on bandwidth allocation. The new methodology utilizes a new Performance Adjustment Factor ("PAF") to allocate bandwidth between GCI's three TERRA service classes.

- Revised TERRA Bandwidth Utilization Summary that uses the new Performance Adjustment Factor ("PAF") to allocate bandwidth between GCI's three TERRA service classes.

- Revised TERRA RoR Cost Study based on a bandwidth allocation to allocate costs among RHC, E-rate and retail services. The cost study has been updated to reflect bandwidth allocation based on the new PAF.

- Revised TERRA RoR Cost Study based on a revenue allocation to allocate costs among RHC, E-rate and retail services. The cost study has been updated to use the bandwidth allocation based on the new PAF to calculate revenue for retail services (*i.e.*, GCI-B and GCI Core bandwidth x [redacted] (25-year per/Mbps) x 12 (months)).

In addition, GCI responds to a question raised by the Division. The Division noted the Brattle report stated that the total capacity on TERRA is [redacted] Mbps, while GCI reported sold capacity of around [redacted] Mbps. Based on this information, the Division asked why the service is oversubscribed when there is all this excess capacity in the network? In supplying capacity data to Brattle, the data was inadvertently not matched to the time periods being compared. TERRA's maximum capacity post-close of the ring in December 2017 was [redacted] Mbps (or [redacted] Gbps), although some segments have lower capacity. Prior to ring close, TERRA's maximum capacity was [redacted] Mbps (or [redacted] Gbps), with the same caveat that some segments had a lower capacity. The corrected capacity measurements do not change Brattle's cross-subsidy conclusion because of the significant difference between Brattle's LRIC calculation and the rate at which GCI imputed revenue for commercial uses to TERRA.

GCI requests confidential treatment of the report submitted herein, as well as the withholding of the designated information from any future public inspection.

In support of this request, GCI hereby states as follows:

1. Identification of Specific Information for Which Confidential Treatment Is Sought (Section 0.459(b)(1))

GCI seeks confidential treatment with respect to the content of this filing, which includes the attachments described above (the "Confidential Information").

2. Description of Circumstances Giving Rise to the Submission (Section 0.459(b)(2))

GCI received information requests from the RHC Telecom Program regarding certain 2017 funding requests of the HCPs for which GCI is a service provider. GCI provided confidential responses to the information requests in November and December 2017 and again on March 30, 2018.[1] Subsequently, GCI met with USAC and FCC staff to discuss the submissions, and the Division has requested that GCI respond to certain proposals and requests regarding the RHC Telecom Program review.

3. Explanation of the Degree to Which the Information Is Commercial or Financial, or Contains a Trade Secret or Is Privileged (Section 0.459(b)(3))

The information for which GCI seeks confidential treatment contains sensitive "trade secrets or privileged or confidential commercial, financial or technical data," which would customarily be guarded from competitors. This is sensitive commercial information that GCI does not otherwise make publicly available. As explained below, public disclosure of these measures could cause competitive commercial harm to GCI. In addition, the mere fact that GCI is being asked to respond may cause competitive harm. Therefore, the information in GCI's

---

[1] *See, e.g.*, Letter from Jennifer P. Bagg, Counsel, GCI Commc'n Corp., to RHC Review, Rural Health Care Program, Universal Serv. Admin. Co. (filed Mar. 30, 2018) ("March 30 Letter").

response constitutes sensitive commercial information "which would customarily be guarded from competitors."

4. Explanation of the Degree to Which the Information Concerns a Service that Is Subject to Competition (Section 0.459(b)(4))

The submitted information contains information regarding GCI's Alaska-based telecommunications services. The Alaskan wireline, wireless, and broadband market (including Ethernet) is subject to competition. In particular, the FCC recently found in the Business Data Services proceeding that the market for Ethernet services is highly competitive.

5. Explanation of How Disclosure of the Information Could Result in Substantial Competitive Harm (Section 0.459(b)(5))

Disclosure of GCI's Confidential Information would cause substantial competitive harm. *First*, disclosure would reveal information regarding GCI's services, including performance characteristics and pricing, and HCP and E-rate customer information. GCI's competitors and customers could use this information to determine GCI's competitive position and associated revenues and thereby gain a competitive advantage. *Second*, disclosure of GCI's Confidential Information would place GCI at a competitive disadvantage, as GCI lacks the same information regarding its competitors. *Third*, disclosure of this information could harm the competitive bidding process in the RHC program.

6. Identification of Any Measures Taken to Prevent Unauthorized Disclosure (Section 0.459(b)(6))

GCI does not distribute the Confidential Information to the public, government officials, competitors, or customers. Each page of the documentation containing any of the Confidential Information is clearly marked in bold-face type "GCI Proprietary – Not for Public Disclosure."

7. Identification of Whether the Information Is Available to the Public and the Extent of Any Previous Disclosure of the Information to Third Parties (Section 0.459(b)(7))

GCI's Confidential Information is and shall remain unavailable to the public. As noted in Part 6 above, GCI has not previously disclosed to third parties, other than the undersigned counsel, any of the Confidential Information.

8. Justification of Period During Which the Submitting Party Asserts that Material Should Not Be Available for Public Disclosure (Section 0.459(b)(8))

GCI requests that the Confidential Information not be disclosed for 10 years from the date of this request. By that time, the sensitivity of GCI's commercial information will have diminished, as market changes will render it increasingly dated, and would make it difficult for competitors to gauge GCI's current market position and revenues.

* * * *

Should you have further questions or require additional information in order to grant the requested confidentiality treatment, please contact me immediately so that I can provide further assistance to resolve this matter.

Sincerely,

JBagg

Jennifer P. Bagg
Counsel to GCI Communication Corp.

Enclosures

cc: Preston Wise

Bandwidth Allocation Methodology to Reflect the Cost-Causative Impact of TERRA Service Class Design and Management

The TERRA network carries multiple classes (or grades) of traffic. The network is designed and managed so that each grade of traffic has different associated service performance parameters. Those parameters will affect the throughput achieved by each class of traffic, as well as the extent to which statistical multiplexing can be used to share capacity among multiple users. In turn, those factors affect how each class of traffic affects decisions to upgrade capacity–and, thus, the extent to which each class "causes" capacity expansion. This white paper outlines the rationale and methodology for calculating an allocation factor based on quality adjusted bandwidth levels.

1. TERRA Network Performance Background

The TERRA network is a middle mile hybrid terrestrial fiber-optic and microwave broadband network connecting Anchorage with other communities in Alaska. GCI provides broadband transport services over the TERRA network including IP/MPLS and Layer Two Ethernet over MPLS service. The longest end-to-end path on the TERRA network traverses [redacted] routers that handle the traffic as it moves from point A to point Z within this middle mile network.

The network transport services delivered by TERRA are used by GCI's customers for a variety of applications. These applications include high performance medical imaging and video conferencing, virtual school classroom extension, wireless service, consumer cable modem service, and commercial internet service. In order to provide a network that can deliver connectivity that satisfies customer performance needs at suitable price points, GCI designed and implemented TERRA with a hierarchical MPLS network architecture. This architecture allows individual packets within the network to be labeled, prioritized and queued according to a set of classification criteria that are defined within the products offered to our customers. The classification consists of three distinct categories; Priority, Normal, and Best Effort. The performance of each service class on the TERRA network is defined by the level of performance by which that service class is managed.

Priority class traffic is GCI's premium TERRA product and is managed to have the lowest latency and packet loss while maintaining superior service availability. In addition, Priority class services are delivered on a one-for-one sold versus provisioned basis. All customers who purchase Priority class service have access at all times to the entirety of the bandwidth purchased. Priority class traffic is placed ahead of Normal or Best Effort class traffic when queuing occurs at any router within the TERRA network.

Normal class service balances price and performance to deliver a better value to customers that may not need the same service performance that is provided at the Priority level. Some of GCI's customers that purchase Normal class service are school districts, state and federal government agencies [redacted] and service providers [redacted]). The service is managed to have a lower service performance and can be statistically multiplexed, yielding a higher efficiency within the pool of available capacity than if all sold Normal capacity were reserved for use of the customer at any time, as is the case for Priority traffic. Normal class service is managed to have lower availability, latency, and packet loss performance than Priority traffic, but better than Best Effort class traffic.

Best Effort class traffic is designed to deliver acceptable performance for consumer-grade Internet applications like web browsing, streaming video (Netflix, Amazon video, Hulu, etc.), instant messaging, and file transfer. Managed to a have lower service performance than Normal traffic, Best Effort capacity is statistically multiplexed at a higher level than Normal class. This traffic is managed to maintain acceptable packet loss and latency performance so that the applications important to customers continue to function. GCI does not provide restoration for Best Effort class traffic in the case of a microwave or fiber network failure event.

Priority and Normal class traffic are the primary drivers of network cost and infrastructure upgrade cycles within the TERRA network. These traffic classes place the highest demand on network resources and require careful engineering to ensure that the performance targets can be met in the operational system. Although upgrades within the network primarily are driven by and support the growth and expansion of Priority and Normal class services, the upgrades also benefit the Best Effort class services due to stair-step increases in network capacity that result from technology upgrades. It is unusual that an upgrade to a network such as TERRA can be tailored or sized to meet only incremental demands. Rather, upgrades are often wholesale equipment or technology replacements and are the result of long-term planning and design to ensure that multiple years of growth can be accommodated within the upgrade project.

Multiprotocol Label Switching (MPLS) is a protocol-independent technology that transports packets and manages traffic between routing points on a network. Measuring the performance of a MPLS network requires the measurement of the quality of the data that is traversing the network on an end-to-end basis within the network (here, the TERRA network). Packets of data can traverse a network in multiple ways. Almost all applications that are transmitted across any IP network today, including MPLS networks, use either Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) as their transport layer protocol. TCP is the most commonly used transport layer protocol because it provides reliable communication, flow control, and congestion control at the transport layer. For instance, the World Wide Web, email, remote administration, and file transfer all utilize TCP. TCP is a stateful/connection-oriented protocol, which means that a handshake must occur in order to facilitate end-to-end communications. Once a connection is set up, user data may be sent bi-directionally over the connection. These connections are also referred to as "TCP sessions" and are managed by end systems connected to the network. Assessing TERRA performance based on TCP permits evaluation of the TERRA network on an end-to-end network basis with respect to what is going to be usable to the applications.

Other traffic on the network that does not require the same level of reliability that TCP provides generally uses UDP. UDP is a stateless/connectionless, rather than a stateful/connection-oriented, protocol. Because it is a connectionless protocol, a dedicated end-to-end connection does not need to be set up. As a result, UDP is less reliable than TCP, but is lighter weight (i.e., requires a smaller packet payload) and has lower protocol overhead. UDP is well suited to real time applications like video conferencing and voice over an IP network.

a. TCP Traffic Performance Provides an Accurate Gauge of TERRA Network Performance

TCP is used in the allocation model presented here as a performance index that represents the class-based service performance differentiation that exists between the Priority, Normal, and Best Effort classes of service. TCP is well-suited as an index because of its inherent latency and packet loss measurements. Specifically, the relative performance among service classes can be established by characterizing the performance of the TCP protocol through the TERRA network. Although not all traffic traversing the TERRA network utilizes TCP, other traffic such as UDP will experience performance levels consistent with the performance for that traffic class. In the case of many UDP streams, understanding the application layer experience is much more subjective because many application layer protocols are proprietary and closed in nature.

[redacted] of traffic on TERRA is TCP[1], and the performance levels to which TCP traffic are managed provide a very accurate, direct gauge on the end-to-end performance of the network. For applications using UDP (or other less common transport layer protocols) the index provided by TCP is a valid measure of the relative performance difference between classes, even though the maximum throughput achieved by UDP may vary from the TCP result. The network performance limitations imposed by packet loss and latency are consistent across all transport layer protocols (TCP / UDP / etc.) and, as such, the relative difference in UDP throughput will be consistent with TCP, but the absolute values achieved may be slightly different.

The use of TCP performance to represent network throughput is a common industry practice as TCP measures round trip time (RTT) to determine how much information can be passed through the network at any instant in time. RTT is also a good measure of how "responsively" an application performs. The lower the round-trip delay, the more responsive the connection feels between the end systems. TCP is sensitive to both latency and packet loss within the network. This sensitivity ultimately limits the TCP session's throughput. Increased latency or packet loss causes the TCP protocol to back off, resulting in lower throughputs and a proportionally degraded experience.

As stated above, the end-to-end performance (as distinguished from availability) of TCP traffic on the TERRA network is primarily dictated by two performance metrics: packet loss and round-trip latency. Packet loss occurs when one or more packets of information traversing the network fail to reach the destination and is measured as a percentage of packets lost with respect to packets sent. Round-trip latency is the time it takes for a packet of information to traverse from end system to another and back and is measured as a unit of time (e.g., milliseconds). In addition to packet loss and latency, overall usability of the service is also impacted by service availability, *i.e.*, the ability to use the service as defined by a service level. Availability is measured as the percentage of time the service is usable with respect to the total time elapsed or, alternatively, the time that a system or network is down relative to a fixed measurement period. Common

---

[1] GCI measured this level at the point where all GCI Internet bound traffic, including for the TERRA network, utilizes the GCI peering connections in Seattle in Portland. The peering routers have monitoring in place to determine the types of inbound and outbound transport layer traffic (which includes TCP, UDP, and any other transport layer protocols). GCI has assumed that TERRA Internet traffic usage is similar to the aggregate traffic breakdown.

measurement periods include monthly and annual availability performance targets (commonly measured in minutes of downtime across those time periods).

b. Performance Is Differentiated by Service Class and Best Measured By the Levels of Service To Which GCI Builds and Manages the TERRA Network

As discussed previously, GCI operates the TERRA network to provide three differentiated service classes—Priority, Normal, and Best Effort. GCI's TERRA network is constructed to provide multiple layers of fault tolerance for its top two classes of services. Priority and Normal class services have access to hardware and microwave path redundancy as well as tertiary satellite backup path restoration in the case of a primary path failure. Although both Priority and Normal traffic classes have access to TERRA's restoration capabilities, Priority traffic has preference when both traffic types are present. Priority traffic will be serviced by the TERRA network, including at each routing point, before Normal class traffic demands are addressed. These network resilience functions significantly increase service availability and allow GCI to offer superior enterprise class service on the network while offering service alternatives with lower performance and a corresponding lower price. The network resilience functions impact the expected packet loss and latency of each service class. Increased packet loss and round-trip latency reduce the theoretical maximum throughput possible for each TCP session and, as a result, the maximum throughput varies by service class.

By specifying the performance metrics for each service class in terms of the latency, packet loss, and availability, an index of the relative performance among the classes can be established. Although packet loss and round-trip time establish service performance under nominal conditions, availability establishes the maximum tolerated downtime.

The performance levels that GCI manages each TERRA service class to are shown below in Table 1.

Class	Availability	Round Trip Latency	Packet Loss	Restoration
Priority Class	[redacted]	[redacted]	[redacted]	[redacted]
Normal Class				
Best Effort Class	[redacted]	[redacted]	[redacted]	

Table 1. TERRA Network Performance Metrics

These performance levels are network-wide. In other words, specific locations on the TERRA network may experience levels that are different from the levels to which GCI builds and manages the whole network.

Assigning an adjustment based on availability, packet loss, and round-trip latency for each class of service is a valid, industry-accepted, standard approach for packet network design.[2] The

[2] Carriers commonly provide tiered service targets based on service class. *See, e,g,, Ethernet Switched E-LAN + Service Level Agreement*, Verizon (2017), http://www.verizonenterprise.com/external/service_guide/reg/cp_eseLAN_plus_sla.pdf.

metrics form the basis for the variation in service classes, and the potential reduction in service performance is tied to an inherent decrease in value associated with each class below Priority. Priority customers value the increased throughput and service availability, which drives greater network investment per purchased bit capacity, and thus pay a premium price for that performance. The purchasing decision is driven by the customer expectation of service performance, and GCI must build and manage the network to support that. Accordingly, these service performance metrics are an appropriate vehicle to measure an approximation of the cost attributable to these services.

The actual instantaneous or bulk utilization of the network (or individual classes) is not sufficiently stable or consistent for allocating costs between services and, therefore, GCI has removed this from the adjustment. GCI regularly upgrades its networks, equipment, and systems to continue expanding capacity to ensure service performance is meeting the established metrics. Although snapshots of service performance can be (and are) captured to ensure that the network is capable of supporting the target service performance metrics, the actual metrics to which GCI designs and manages the network are the cost drivers of the network. The forecasted capacity that any class of service consumes must be considered in tandem with how that capacity drives network investment. Specifically, Priority traffic capacity comes with the lowest round trip latency, the lowest packet loss, and the highest availability. The network can only achieve those performance metrics if the aggregate traffic distribution on the network consists of a mix of other traffic types. That is, every bit of traffic on the network cannot be assigned to Priority class service or the performance levels for that service class will not be met because TERRA is a constrained resource. More demand exists than total capacity of the TERRA network and, as such, some of the network traffic must experience a performance degradation to ensure sufficient capacity for the Priority traffic.

2. Alternative Approach Eliminates "Oversubscription" as an Independent Adjustment Factor

In proposing this approach, GCI has removed the relative utilization and sold capacity measures from the allocation model. The relative weight to assign to utilization versus the level to which GCI designs and manages the network to ensure service performance is difficult to quantify in any but instantaneous analyses. As discussed above, the instantaneous analysis is not practicable to apply over time.

In its initial bandwidth adjustment, GCI included an "oversubscription" adjustment that was based on actual 2017 and 2018 network utilization measurements. The adjustment is more accurately described as a tiered service utilization percentage as it is calculated by dividing the measured utilization by the provisioned bandwidth by service tiers (priority, normal, and best effort). The factor represents only the traffic placed on the network at a specific point in time by class and, therefore, is highly variable, which makes it difficult to provide a true measure over time of sold capacity versus demand placed on the network. In addition, utilization measurements provide a snapshot of utilization at the point of measurement, and do not reliably represent total network utilization. This snapshot may show that a specific location requires an upgrade at that location, but it is not indicative of end-to-end network cost causation.

GCI had attempted to calculate an overall network oversubscription model that included all services (*e.g.*, rural broadband, cable modem, mobile wireless broadband). However, based on the follow-up discussions and subsequent review, it became apparent that this calculation was not sustainable as a standalone or additional measure. Any measurement of network oversubscription is instantaneous in nature and, therefore, highly subjective. Oversubscription is very dependent on customer behavior and use case. For example, a consumer browsing the internet can be oversubscribed many, many times while a clinic transferring x-ray images depends on intense resources for much longer periods of time. The performance of the internet web traffic TCP sessions relative to the x-ray image file transfer, however, is a good measure of service performance. Additionally, by providing dramatically fewer minutes of downtime (40x for Best Effort versus Priority), the service continuity for Priority traffic provides increased value to those customers. The physical port capacity within packet switching devices is frequently designed to be 10x the available link capacity and, therefore, is irrelevant to a network utilization analysis. Port capacity also provides no insight into the performance of the sessions or streams of data flowing through the network, because only the individual packets are being switched in and out of a router.

GCI's further review of the network utilization/oversubscription factor revealed that it does not actually create a cost driver, in that it has little relationship to whether network upgrades are deployed. Therefore, GCI determined that an alternative factor should be developed that would be straightforward to calculate, sustainable over time, and rationally related to cost causation principles.

3. Performance Adjustment Factor (PAF) Analysis

The PAF is a function of the three previously discussed input parameters: packet loss, round trip latency, and availability. These parameters are used to compute a reduced theoretical maximum throughput at the guaranteed performance value.

- The three computed adjustment factors are equally weighted and the PAF is the average value of the three adjustment factors. Averaging is a reasonable approach because it equally weights the importance of packet loss, round trip latency, and availability when evaluating the TCP sessions experience for end user systems. Weighting one factor over another would impose an arbitrary and subjective valuation into the calculation.

- The PAF is defined relative to Priority traffic. That is, Normal and Best Effort classes will show performance reductions against the highest performing Priority class. In other words, Priority traffic is treated as the baseline (i.e., the service with the greatest maximum guarantee throughput), with reductions applying only to Normal and Best Effort at tiered values.

Using work previously done by Mathis, *et. al*.[3] it can be shown that the maximum throughputs under latency and packet loss are limited. The computed theoretical maximum throughputs for TCP sessions under the performance constraints for Priority, Normal, and Best Effort classes of

[3] *See* Matthew Mathis, et. al., *The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm*, Computer Communication Review, July 1997, at 2, eqn 3.

traffic are as shown in Table 2 along with the computed annual downtime for each class based on the stated service level availability guarantee.

An example calculation applying the Mathis equation to compute maximum throughput for the Normal class with packet loss and nominal latency is shown below.

$$B = \frac{MSS}{RTT} * \frac{C}{\sqrt{p}}$$

The default maximum segment size (MSS) for common operating systems is 1460 bytes, the RTT is the latency as defined for each class and p is the packet loss rate. C is a constant of proportionality and is related to various TCP implementation combinations and loss mechanisms. For the purposes of this calculation a value of unity is assumed. See Table 1 in Mathis for more detail.

$$B_{Normal} = \frac{1460 \text{ bytes} * 8 \text{ bits/byte}}{50 * 10^{-3} \text{seconds}} * \frac{1}{\sqrt{0.01}} = 2.3 \text{ Mbps}$$

With regard to service availability, the relative downtime values among the three service tiers is necessary to understand the performance differential and cost impact. Outage duration for priority class service is measured in terms of minutes or hours over a year, while Best Effort service outage produces measurements in terms of days. As shown below, Priority class service is managed to less than [redacted] of total outage over the course of a year (*i.e.*, [redacted]), which results in a [redacted] service availability, and Best Effort class service is managed to around [redacted]—or approximately [redacted]—of total outage over the course of a year (*i.e.*, [redacted]), which results in a [redacted] service availability. Moreover, the difference in cost to provide service availability of [redacted] for Priority class service, versus [redacted] or [redacted] for Normal or Best Effort class service is huge, and each order of magnitude increase in availability is exponentially more expensive. In other words, by communicating in terms of a downtime multiplier the relative impact on service performance of the different classes of services can be shown, as well as the cost impact of service availability on the performance adjustment factor.

Class	**Latency (maximum TCP rate)**	**Packet Loss (maximum TCP rate)**	**Availability (annual downtime)**
Priority Class	[redacted]	[redacted]	[redacted]
Normal Class	[redacted]	[redacted]	[redacted]
Best Effort Class	[redacted]	[redacted]	[redacted]

Table 2. TERRA class-based latency, packet loss, and availability based intermediate results

Using the computed maximum TCP rates for latency and packet loss along with the availability, the individual factors contributing to the overall PAF can be determined. The factors are shown in Table 3 along with the computed PAF value for each traffic class.

Class	Latency Factor	Packet Loss Factor	Availability Factor	PAF
Priority Class	[redacted]	[redacted]	[redacted]	[redacted]
Normal Class	[redacted]	[redacted]	[redacted]	[redacted]
Best Effort Class	[redacted]	[redacted]	[redacted]	[redacted]

Table 3. TERRA class-based latency, packet loss, and availability-based adjustment factors and computed PAF

The PAF is computed by taking the average of the latency and packet loss impact factors. These factors are computed as the percentage reduction observed relative to priority class traffic.

$$PAF = \frac{F_{latency} + F_{packet\ loss} + F_{availability}}{3}$$

For example, the Normal Class PAF value is computed as shown below.

$$F_{latency}(\text{Normal}) = \frac{[redacted]\ \text{Mbps}}{[redacted]\ \text{Mbps}} = [redacted]$$

$$F_{packet\ loss}(\text{Normal}) = \frac{[redacted]\ \text{Mbps}}{[redacted]\ \text{Mbps}} = [redacted]$$

$$F_{availability}(\text{Normal}) = \frac{[redacted]\ \text{min}}{[redacted]\ \text{min}} = [redacted]$$

$$PAF = \frac{F_{latency} + F_{packet\ loss} + F_{availability}}{3} = [redacted]$$

Having determined the per-class PAF, this factor is applied to each class's provisioned service capacity to determine the adjusted capacity associated with that service class in Table 3.

4. Results of the PAF as Applied to the Bandwidth Allocation Model

Applying the results in Table 3 to the TERRA bandwidth allocation model shifts the allocations of capacity between the priority, normal, and best effort classes.

In the previously submitted filing, the RHC, E-Rate, and Other traffic types were distributed across the bandwidth as shown in Table 4. In this table, the alternative allocation, the previously submitted allocation, and the delta between the two are shown.

The primary impact of the alternative approach is that slightly more allocation is shifted from the Best Effort traffic type to the E-Rate and RHC traffic types (which include both Priority and Normal class traffic).

	2014 alt / old / delta	2015 alt / old / delta	2016 alt / old / delta	2017 alt / old / delta
RHC	[redacted]	[redacted]	[redacted]	[redacted]
E-Rate	[redacted]	[redacted]	[redacted]	[redacted]
Other	[redacted]	[redacted]	[redacted]	[redacted]

Table 4. Traffic allocation by customer type for 2014-2017 time period for alternative and previously submitted approaches

Revised TERRA Bandwidth Utilization Redacted in Entirety

Revised TERRA ROR Cost Study (Bandwidth) Redacted in Entirety

Revised TERRA ROR Cost Study (Revenue) Redacted in Entirety